A Multi-criterial Decision Support System for Forest Management

Donald Nute, Geneho Kim, Walter D. Potter

Artificial Intelligence Center
111 Boyd Graduate Research Center
The University of Georgia
Athens, GA 30605
dnute@ai.uga.edu

Mark J. Twery, H. Michael Rauscher, Scott Thomasma, Deborah Bennett, Peter Kollasch

USDA Forest Service

Abstract

We describe a research project that has as its goal development of a full-featured decision support system for managing forested land to satisfy multiple criteria represented as timber, wildlife, water, ecological, and wildlife objectives. The decision process proposed for what was originally conceived of as a Northeast Decision Model (NED) includes data acquisition, goal selection, goal satisfaction analysis, goal conflict analysis, stand modification recommendations, silvicultural method assignment, and stand treatment recommendations. NED-1 supports the first three functions. We discuss the AI techniques that are used in NED-1 and that will be used to add the remaining functions planned for later versions of NED.

Introduction

NED is an acronym derived from "Northeast Decision Model". Originally intended for use in managing national forests in the Northeastern United States, NED has evolved into a decision support tool for managing both public and private forested land throughout the eastern United States. Unlike many forest management decision support systems that have timber production as their single objective ((Lorenzo 1993), (Nute et al. 1995a), and (Nute et al. 1995b) for example,) NED is designed to help managers plan for wildlife, ecology, water, and landscape objectives as well as timber production. A description of NED from the environmental scientist's point of view can be found in (Twery et al. 1997) and (Twery et al. 1998). In this paper, we will describe the NED project from the decision support system developer's perspective with emphasis on the AI techniques involved.

The NED project has as its objective the development of a full-featured decision support system for managing forested land to meet multiple criteria or multiple objectives. An initial version of NED has been released, one which achieves only part of the function we hope eventually to capture in NED. This version represents at least five years of intensive effort by a team including three full-time programmers, two artificial intelligence researchers, and more than a dozen domain experts in silviculture, ecology, wildlife management, and other disciplines important for forest management. First, we will describe a decision process for forest management that we hope eventually to integrate into NED. Second, we will describe how NED-1, the first release of NED, functions. Third, we will discuss the tasks that remain concentrating on those tasks involving AI techniques.

A Decision Process for Forest Management

A management unit is a piece of forested land divided into several stands where each stand is a contiguous piece of land representing a single forest type. A management unit might include 50 acres or 50,000 acres and might contain half a dozen stands or several hundred stands. A management plan or a forest includes overall goals for the entire management unit plus decisions about how to manage each stand in the unit. The management plans for the individual stands should together accomplish the overall goals for the management unit. The management plan is then used to produce short-term treatment recommendations for individual stands in the forest.

The decision process for NED will be non-deterministic. First, the process will be non-linear and iterative. There is no single order in which steps in the process must be taken, and it is normal to return to earlier steps at different points in the process. Second, the process is not expected to produce one management plan even for a single management unit and a single set of objectives or goals. Planners and managers may reject all or parts of the plan or recommendations produced by NED and force it to generate alternative plans or recommendations. Planning and management must be continuous and flexible because goals change over time, because managers have their individual preferences about how to manage forested land, and because forests experience natural disturbances beyond the manager's control or ability to foresee. Nevertheless, the management process involves certain kinds of activities and some of these do have a natural order or tend to be revisited at certain points in the ongoing management process.

The first step in the proposed decision process is data acquisition: we must provide a description of the management unit. We need to know the location and number of stands in the forest. For each stand we need information about size, shape, species and size classes of trees in the overstory, overstory closure, understory density, presence or absence of water, etc. We also need information about how the stands are related to each other spatially and about the accessibility of the stands.

In addition to stand information, we must identify management objectives for the management unit in a goal selection phase. These can include timber, wildlife, water, ecological, landscape, and recreational objectives. A possible wildlife objective would be to manage the forest to provide habitat for a particular species, a possible water objective would be to protect an existing watershed, and a possible landscape objective would be to maintain a certain visual quality. Some ecological objectives such as maintaining the health of the trees in the forest are assumed, but other possible ecological objectives would include maintaining a certain degree of biological diversity or maintaining or improving soil nutrients.

Goal satisfaction analysis can be performed once management objectives and information about current conditions in the forest have been. The purpose of this analysis is to determine how well the forest currently satisfies the decision-maker's objectives. Goal satisfaction analysis can also be performed later in the recommendation phase of the decision process. As different scenarios for the forest are generated, goal satisfaction analysis can be performed on each scenario to determine how well it will satisfy the decision-maker's objectives.

Besides goal satisfaction analysis, another important step in the decision process is goal conflict analysis. The purpose of this step is to determine whether it is probable or even possible that all of the decision-maker's objectives can be satisfied given the constraints of the management unit. It might be impossible to produce a given amount of timber from a small forest or to maintain a certain species of wildlife in certain regions of the country. In these cases, the objectives conflict with constraints placed on us by the forest we are trying to manage. In other cases, the conflicts may arise out of the objectives themselves without regard to the management unit. To manage for a particular large wildlife species could require most of the forest to be of a certain type, while timber objectives could require most of the forest to be of a completely different type. While it would be possible to manage the forest to produce either of the two types, it would obviously be impossible to manage it in a way that produces two different types over most of the management unit.

Any conflicting goals that are discovered should be resolved before proceeding to later phases of the decision process. Of course, goal conflicts could also be discovered at later stages in the process, forcing reconsideration of the goal set. An extended discussion of the knowledge representation problems associated with developing a goal structure that will support goal satisfaction analysis, goal conflict analysis, and goal conflict resolution can be found in (Nute et al. 1999)

Assuming an adequate description of the management unit and a feasible set of management goals or objectives, the next will be to arrive at general recommendations about how to manage the forest. The previous goal could indicate that the general structure of the forest is already appropriate for the objectives. Then the recommendation would be to maintain this structure. However, the structure of the forest may have to be modified to maintain a wildlife species, to protect a watershed, to reduce threats to forest health, to improve timber production, to improve the aesthetic and recreational value of the forest, or to accomplish any combination of these and other objectives. Assuming that the stand structure for the forest will remain constant, these recommendations will take the form of suggested modifications for the structure of individual stands. One way to approach this is to find the smallest modification to the smallest number of stands that will allow satisfaction of all the objectives that have been established for the forest. Various modeling tools will be needed both to evaluate how well a proposed structure will satisfy the objectives and to determine the best methods for converting the target stands. Managing these modeling tools becomes a central problem for developing NED. The form of the general recommendation will be a list of stands to modify together with suggestions about how to accomplish the modifications.

A silviculture method will be assigned to each stand that already has the desired structure. These silviculture methods can then be used to generate specific treatment recommendations for these stands.

These treatments will include recommendations about when and how to regenerate, thin, fertilize, and harvest the stand. In some cases, assigning a silviculture method may also be enough to tell us how to modify a stand whose must be changed to satisfy the overall objectives for the forest. If an open stand needs to be converted to an even-aged single species stand for timber production, the appropriate silviculture method should recommend both that the stand needs to be regenerated and that the method for regenerating it.

The decision process described here proceeds from a description of the forested land to be managed and an initial set of management objectives, to general recommendations about how to modify the forest and specific treatment recommendations for individual stands within the forest. To implement this decision process in a decision support system, we will need a combination of sophisticated interface tools, knowledge based systems, modeling tools, and an intelligent model management system. We also envision the possible use of artificial neural nets and genetic algorithms at certain points in the process. We will say more about this after a quick look at the current status of NED.

NED-1

The first version of NED was released in March of 1999, and the U.S.D.A. Forest Service has conducted several training programs for potential users. NED-1 is an *analytic* version of NED that implements the proposed decision method through the goal satisfaction analysis stage.

Figure 1 combines a functional diagram of NED-1 with a representation of the current NED architecture. NED includes both declarative and procedural components. Our interest here is in the declarative components implemented in DSSTools, a Prolog toolkit for Decision Support System development created at the Artificial Intelligence Center of the University of Georgia through collaboration with the Southeast Forest Experiment Station and the Northeast Forest Experiment station of the USDA Forest Service. In this paper, we will focus on DSSTools and the AI components of NED-1.

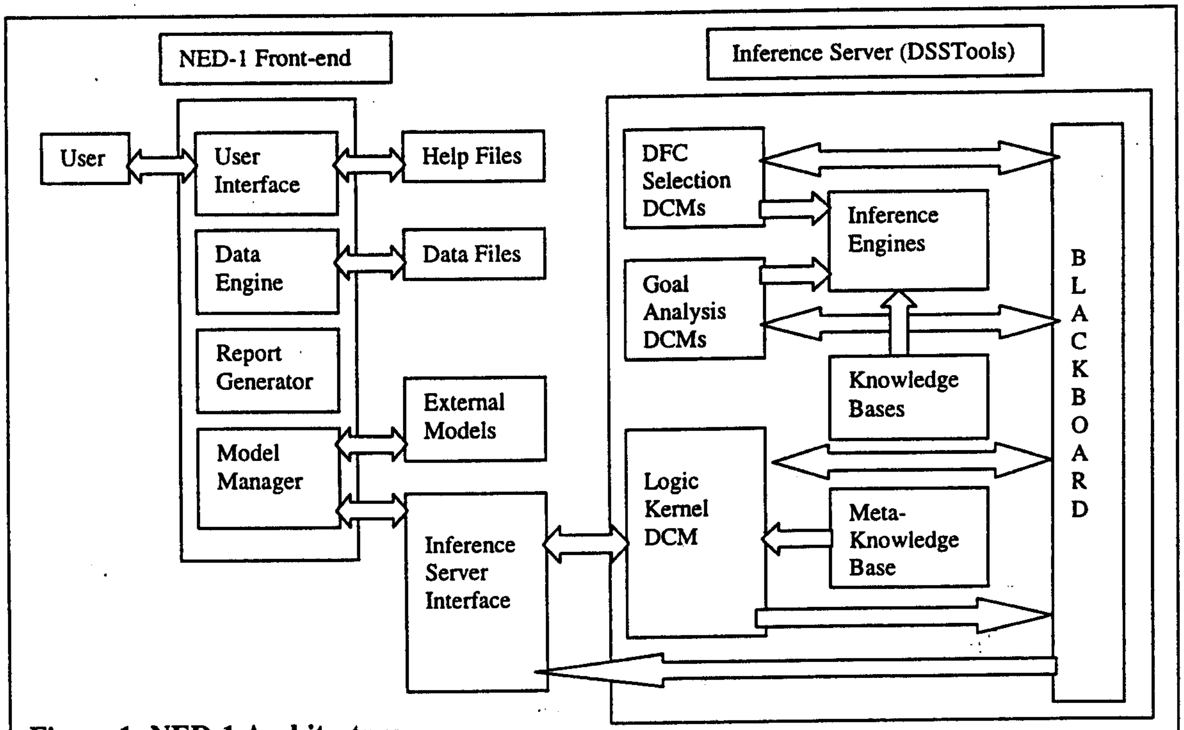

Figure 1: NED-1 Architecture

An inference server DSSTools application consists primarily of one or more knowledge bases and one or more *domain control modules* (DCMs.) The DCMs are semi-autonomous agents that have access to the blackboard and can call any of the inference engines available in DSSTools. Only one DCM is active at any time. After a DCM terminates, a scheduler gives each DCM an opportunity to examine the blackboard. The first DCM whose activation conditions are satisfied by requests and information on the blackboard then becomes active. Several DCMs can play a role in responding to a single request for inference, each completing its piece of the process. But DCMs do not communicate their results to each other directly. Instead, they write their results to the blackboard where every DCM can see them. A DCM "knows" when it should perform its task by looking at the facts on the blackboard, including requests for specific services that that DCM may provide. Similarly, when a DCM calls an inference engine, the inference engine does not report its results back to the DCM that called it. It writes its results to the blackboard where they become available to every DCM.

There are three kinds of DCM included in NED-1. Most of these are paired with a set of facts and rules that provide the knowledge to perform some essential function within NED-1. Each DCM, together with its associated knowledge base and whatever inference engine it uses, constitutes a small knowledge based system within NED-1.

The Logic Kernel DCM receives requests from the NED-1 front-end program. For example, the Inference Server could be requested to determine the dfcs for the wildlife goals that have been selected or to provide a list of wildlife species for which suitable habitat can be found on the management unit. The NED-1 front-end program has all the data concerning the management unit, but it does not know what data will be required to fulfill any of these requests. The Logic Kernel DCM consults its meta-knowledge base to determine which knowledge base contains the facts and rules needed to fulfill a request, loads the knowledge base, and puts the request on the blackboard. The appropriate DCM sees the request and performs the necessary inference. While the DSSTools component of NED-1 knows how to solve various problems, it has no direct access to the forest data. It requests this information from the NED-1 front-end program as needed. When inference is complete, the NED-1 front-end is informed. It can then call the Inference Server Interface to read the results from the blackboard. When the Inference Server receives a new request, the Logic Kernel DCM erases the blackboard so what may be obsolete conclusions will not affect subsequent inference.

NED-1 has a simple two-level goal structure. The user selects management objectives from menus. Each of the higher-level timber, wildlife, and water objectives is associated with a set of lower-level *desirable future conditions* (dfcs) using a knowledge base specific to the category of the objective. A dfc is an observable variable together with a desired value for that variable. (See (Twery et al. 1998) and (Nute et al. 1999) for more details.) The other two kinds of DCM in NED-1 deduce sets of dfcs from goals or evaluate the forest to see how well it satisfies a set of dfcs.

The rules for evaluating how well the forest satisfies user objectives are "fuzzy". The idea behind fuzzy set theory is that an object may belong to a class "more or less". For example, a $10,000 car definitely does not belong to the class of expensive cars and a $50,000 definitely does. But what about a $25,000 car? The price range for expensive cars is fuzzy at its lower end. The values associated with the observable variables in dfcs represent thresholds that can have this kind of fuzziness. In NED-1, any value that is 5% above the threshold in the dfc is considered to clearly satisfy the dfc, and any value 5% below the threshold is considered to clearly fail the dfc. Values within 5% of the threshold are reported as marginally satisfying or marginally failing the dfc.

While most NED-1 objectives are handled by knowledge based systems as described, landscape and ecology objectives are handled differently. A Prolog routine uses management unit information together with a stand adjacency table to perform a "patch analysis" on the forest using forest type, size class, or a combination of the two to determine the patches. Different methods for evaluating the amount of "patchiness" can then be applied to this analysis.

The ecological component of NED-1 is a Forest Health system that can identify associated causes for problems in overstory species. This system includes 118 insects, viruses, funguses, wildlife, human, and non-biotic agents that are associated with damage to 45 different overstory species. After receiving information from the NED-1 front-end program about the composition of the overstory, Forest Health takes user information about tree damage and reaches conclusions about the likelihood that the damage is caused by any of the agents included in the system. Originally written in a procedural language, the Forest Health system is included as an external model in NED-1. It is being converted to a knowledge based system in Prolog. This will make it easier to maintain and to expand the system.

Future NED Development

The next version of NED will not only add to the functionality of the system; it will also have a radically different architecture. One of the weaknesses of the NED-1 architecture is the procedural front-end program handles model management. In the next version, by moving this crucial function to the DSSTools part of the system will enable truly intelligent model management. Current plans are for top-level control of future versions of NED also to be provided by the DSSTools. Figure 2 is a simplified diagram of the architecture for NED-2 and later versions of NED. We begin by looking at some planned changes in the procedural components of NED that will make it easier to move overall control of NED to DSSTools, Data entry and other dialogs incorporated into the NED-1 front-end program will be separated into distinct procedural modules (Windows dynamic link libraries or DLLs) to facilitate modular development. These modules will be called by DCMs. New user interface dialogs will be implemented with user interface tools in DSSTools or as additional external modules. The custom database system used in NED-1 will be replaced by an MS Access database accessible to the user interface modules and all DCMs in the system through the Data Manager and the blackboard. The help files and forest management documents in NED-1 will be converted to HTML files and accessed using a Web browser. In subsequent versions of NED, reports will be generated as HTML files by new DCMs. A rule-based system for generating reports will improve the quality and flexibility of the reports. DCMs for generating the reports found in NED-1 have already been developed. New DCMs will be required to generate additional reports as new functionality is added to NED.

Meta-knowledge about how to use different models and visualization tools will be developed. Tools already in use with the USDA Forest Service, including the Forest Vegetative System (FVS) and the Stand Visualization System (SVS), will be integrated into NED by building "wrappers" that permit communication between a model and a DCM. Besides these and other "legacy" or "third-party" models, new models specifically designed to plug into NED will be developed including a soil nitrogen model. The meta-knowledge will allow NED to respond to requests like, "Show me how Stand 46 will look in 25 years if I remove all hardwoods under six inches in diameter today." Given this request, a future version of NED will create a temporary database for the modified stand, run an appropriate growth and yield model on the modified stand data, and send the results of the growth and yield model to a visualization routine. The knowledge for performing such complex tasks will be stored in declarative meta-knowledge bases. Initial work on integrating FVS, SVS, and other important existing models and visualization tools is nearly complete.

Modeling regeneration of disturbed stands has proven to be a particularly difficult problem. One approach using a logical model based on (Loftis 1990) is nearly completed. Another empirical approach using an artificial neural net is reported in (Berkshire 1995). The logical model can be incorporated into a DSSTools DCM. Another tool in DSSTools allows an application to load a trained neural net created using the commercial NeuroShell package and feed inputs to it. This tool pulls data from the blackboard, feeds it to the trained neural net, and puts the outputs of the neural net on the blackboard where all DCMs can see it. Both of these approaches to regeneration will be explored for later versions of NED.

NED-1 requires the user to provide detailed information about the management unit including tree inventories for representative plots for each stand in the forest. Collecting this information is too expensive for the National Forests and other large tracts of forested land. Different data imputation methods will be explored using aerial photography and a library of profiles for forested regions in the eastern United States. This system may use neural nets, case-based reasoning, or a combination of these and other AI techniques.

A better model of the goal structure for NED must be developed before goal conflict analysis will be possible. We also need a model of the causal and other relations between the dfcs included in the goal structure since most of the conflict analysis will be carried out at this level. Most of the goal conflict analysis will be performed by a knowledge based system, but the use of growth and yield models, wildlife models, and other simulations. May augment this.

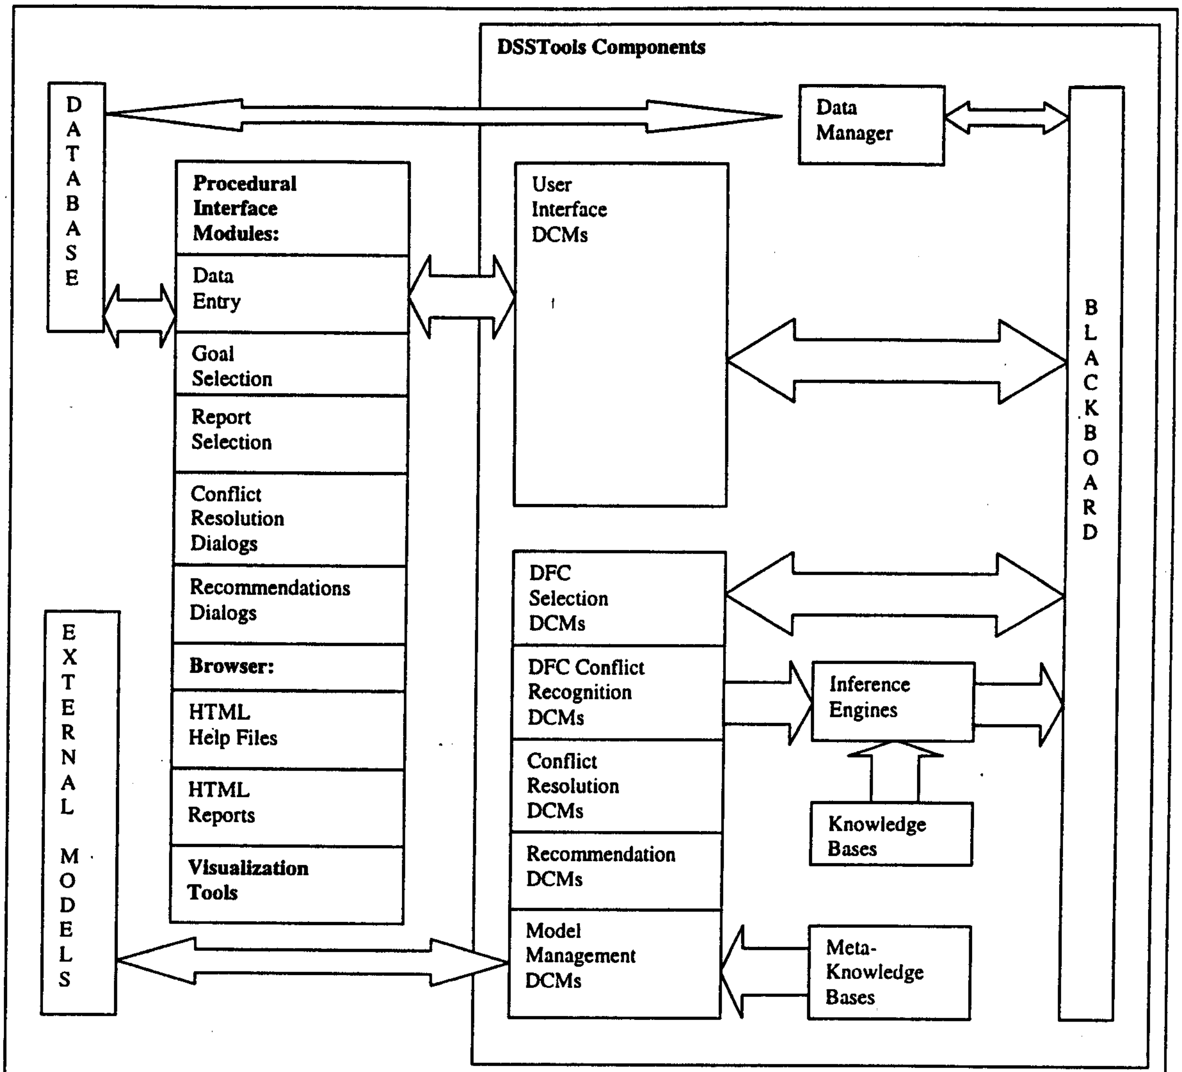

Figure 2: Future NED Architecture

Knowledge bases to support general recommendations in NED will be developed. NED–1 provides no recommendations other than those incorporated into the text of the hypertext forest management and help documents. A first step will be to explain to the user how a particular objective fails. For example, NED-1 can determine whether the forest provides habitat for a particular wildlife species. It can also provide a list of all species for which suitable habitat can be found in the forest. But the knowledge base is not developed in a way that allows NED-1 to tell the user what features are missing if the forest does not provide habitat for a species. The same is true for the other categories of objectives. A first step toward general recommendations, a step beyond the simple goal satisfaction analysis now provided, will be a more detailed analysis of exactly how and why the management unit fails to meet various objectives.

A further step toward a version of NED that can make general recommendations will be a knowledge based system that can find a plausible set of modifications to the stands in the management unit that would produce a situation in which all objectives can be satisfied. Part of this task will depend on specifying which stands will be used to satisfy which management objectives. Of course, we want to minimize the number of stands that will have to be modified, and we will want to minimize the scope of the modification for each target stand. This constraint derives from our assumption that resources for managing the forest are also constrained. A knowledge based system in a future version of NED will determine what kinds of stands are needed to satisfy failed objectives. Then it will identify stands most similar to the target stands. Next the future NED will substitute the projected modified stands for the existing stands and use simulations to predict how the non-target stands will have changed during the time it will take to modify the target stands. Finally, something like the knowledge based system now in NED-1 will perform goal satisfaction analysis on the projected state of the forest. This generate-and-test procedure will be the most computationally intensive part of the NED system. Besides complex knowledge bases, it will require intelligent model management to set up and execute a wide range of external simulations.

Once a suitable scenario for modifying the forest has been proposed and accepted, rule-bases will be needed for assigning appropriate silviculture methods to each of the stands in the forest. While research will continue to produce changes in the ways we manage certain kinds of stands, there already exists a good practical understanding of many of the relevant issues. These methods will have to be incorporated into knowledge bases. Examples include the knowledge bases developed for even-aged stands of red pine (Nute et al. 1995a) and aspen (Nute et al. 1995b) and a knowledge base developed for uneven-aged stands of loblolly and shortleaf pine (Lorenzo 1993). Specific objectives must also be assigned to various stands. With current stand information, a set of objectives for the stand, and a knowledge base capturing the assigned silvicultural method, NED will be able to make specific treatment recommendations for particular stands.

Conclusion

We have described a decision process for managing forested land to satisfy multiple criteria or achieve multiple objectives. NED-1 is a software system that helps a user perform some of the steps in this process. NED-1 uses a blackboard architecture and knowledge based systems for the inferential portion of the decision process. But in this initial implementation the benefits of these AI techniques do not extend to overall control of the system or to the management of the many simulations and other models that are needed for multi-criterial forest management.

A revised architecture for NED has been developed that will extend the blackboard architecture to the entire system rather than just the inference server. Using this new architecture, future versions of NED will coordinate rule-based systems, mathematical simulations, neural nets, and visualization tools in an intelligent manner. Meta-knowledge will be used to coordinate these disparate components in a way that will be transparent to the user.

While a particular decision model guides its development, NED is intended to be a flexible decision support system that allows users to employ the different functions supported by NED in a variety of ways. There is no definite order in which the user must use the different tools contained in NED. This flexibility is compatible with the iterative, non-linear decision model we are using. Forest management takes place in a world where natural events, availability of resources, and the desires of society or of the individual forest manager change. A useful tool to help manage a forest or any other ecological system should have enough flexibility to allow the user to explore different possibilities to find solutions that are most comfortable and to modify previous solutions under changing conditions.

References

Berkshire, John. 1995. *Simulating Allegheny Understory Regeneration and Growth Using Artificial Neural Network Modeling Techniques.* Masters thesis, Artificial Intelligence Center, The University of Georgia, Athens, Georgia.

Liu, S.; Potter, W. D.; and Rauscher, H. M. 1999. Using DCOM to Support Interoperability in Forest Ecosystem Management Decision Support Systems. International Union of Forestry Research Organizations Conference: The Application of Scientific Knowledge to Decision-making in Managing Forest Ecosystems, Asheville, North Carolina, pp. 3-7.

Loftis, David L. 1990. Regeneration of Southern Hardwoods: Some Ecological Concepts. In *Proceedings of the National Silvicultural Workshop, July 10-13, 1989,. Petersburg, Alabama.* USDA Forest Service, Washington, DC., pp. 139-143.

Lorenzo, Alfredo B. 1993. *Development and Evaluation of an Expert System Approach to Uneven-aged Management of loblolly-Shortleaf Pine Stands in the West Gulf Region.* Ph. D. dissertation, School of Forestry, Wildlife, and Fisheries, Louisiana State University, Baton Rouge, Louisiana.

Nute, Donald; Rauscher, H. Michael; Perala, Donald A.; Zhu, Guojun; Chang, Yousong; and Host, George M. 1995a. A toolkit approach to developing forest management advisory systems in Prolog. *AI Applications* 9(3):39-58.

Nute, Donald, Rauscher, H. Michael; Perala, Donald A.; Zhu, Guojun; Chang, Yousong; Host, George M.; and Worth, Christopher V. 1995b. An aspen forest management decision support system implemented using a Prolog knowledge systems toolkit. (Abstract) In J. Michael Power, Murray Strome, and Terry C. Daniels (eds.), *Proceedings of Decision Support - 2001*, American Society for Photogrammetry and Remote Sensing, Bethesda, Maryland, p. 686.

Nute, Donald, Nath, Shree; Rosenberg, Gregg; Verma , Brahm; Rauscher, H. Michael; Twery, Mark J.; and Grove, Morgan. 1999. Goals and a Goal Orientation in Decision Support Systems for Ecosystem Management. International Union of Forestry Research Organizations Conference: The Application of Scientific Knowledge to Decision-making in Managing Forest Ecosystems, Asheville, NC, 3-7, 1999.

Twery, Mark J., Bennett, Deborah; Kollasch, Peter; Thomasma, Scott; Stout, Susan; deCalesta, David; Hornbeck, Jim; Steinman, James; Miller, Gary; Grove, Morgan; Rauscher, H. Michael; Gustafson, Eric; Cleveland, Helene; Palmer, James; Hoffman, Robin; McGuinness, Barbara; Chen, Ningyu; and Nute, Donald. 1997. NED-1: an integrated decision support system for ecosystem management. In *1997 ACSM/ASPRS: Annual Convention and Exposition Technical Papers*, Vol. IV, American Society for Photogrammetry and Remote Sensing and American Congress on Surveying and Mapping, Bethesda, Maryland, pp. 331-342.

Twery, M. J., Stout, S. L.; and Loftis, D. L.1998. Using desired future conditions to integrate multiple resource prescriptions: the Northeast decision model. In S. A. El-Swaify and D.S. Yakowitz (eds.), *Multiple objective decision making for land, water, and environmental management, proceedings of the first international conference on multiple objective decision support systems for land, water, and environmental management: concepts, approaches, and applications,* 23-27 July, 1995, Honolulu, HI. CRC Press LLC, Boca Raton, FL, pp. 197-203.